EXHIBIT B

Explainer FAQ for TURTLEDOVE

Explainer: go/turtledove internally, https://github.com/michaelkleber/turtledove externally

What use case is TURTLEDOVE trying to meet?

This is an API for allowing ads targeted at an advertiser-created audience, to replace targeting a campaign at a list of cookies. The fact that a particular person is in the target audience is only known to that person's browser, and a person can only be added to an audience by running javascript code in their browser. Showing ads to the people in the audience would involve ad tech adopting a new in-browser auction.

So is this a way to do remarketing without cookies or other tracking?

Yes.

What limits would the browser place on audiences?

We don't know yet. Browsers would be in a position to impose limits like "audience membership can only last for 30 days" or "audiences must have at least 100 people in them". But this is just an Explainer, the beginning of the discussion of an idea. Any limits would be the topic of public discussion.

How is this related to PIGIN?

PIGIN also proposed browsers holding audience memberships, but showing ads to people in the audience involved the browser revealing audience membership to the server as part of the normal ad request. That didn't do a good enough job of keeping audience memberships truly private. TURTLEDOVE replaces PIGIN, and fixes the privacy problem using an in-browser auction.

Is the browser picking which ad wins?

NO! The winning ad is chosen by the publisher's ad network, just like today. The difference is that some of the ad network code that chooses the winning ad is Javascript running inside the user's browser.

What about RTB, exchange buying, etc?

TURTLEDOVE can support advertisers using multiple DSPs, all of which can buy on multiple SSPs. The Explainer is vague on the details here, and it will require new server-to-server protocols, so there are a lot of details that will need to be worked out — in a discussion invoking both browser makers and the various ad networks that want to work with each other.

What is the timeline for this change?

The Explainer has a section called "Incremental Adoption Path" that describes a bunch of possible initial steps along the way to the full privacy-preserving way of showing ads. There is no timeline attached to any of them. The first (Minimum Viable Product) step has fewer security guarantees, but is already enough to support remarketing ads without cookies, and this step will be in use before 3rd-party cookies go away.

Will there be any new restrictions on what ads can be shown using TURTLEDOVE?

By the later stages of TURTLEDOVE, the audience-targeted ads will need to be a new sort of creative that can be loaded in advance and then rendered in a privacy-preserving way, without further connecting to the network or communicating with the surrounding page. This new work has a longer time horizon, and won't be required for the MVP.

Also in the later stages, audience-targeted ads will need to use new privacy-preserving aggregated measurement tools for their reporting needs. This likewise has a longer time horizon, and won't be required for the MVP.

Is it going to be a lot of work to adapt to this new way of doing things?

Yes. Sorry.

This is just an Explainer, so a lot may still change in the ensuing discussion. But what's proposed now is the simplest-to-adopt thing we've been able to come up with that has the necessary privacy properties. Remarketing without tracking is unavoidably new and different, and will be some serious work, no matter how we slice it.